Trivial Baselines as Redundant Measurements

R. Cliff Wilkie

ABSTRACT: Since GPS technology is a relatively new measuring technique, methods for achieving redundant measurements have not been clearly established and universally accepted. In particular, dependent or "trivial" baselines are a recurrent source of disagreement. Some GPS users consider them redundant measurements, and some do not. In order to evaluate the issues involved, an examination of the fundamentals of the GPS measuring process is made by analogy to conventional surveying methods. What constitutes the actual measurements in GPS phase differencing methods is shown as well as how these actual measurements relate to the GPS vectors that are usually considered as measurements in network adjustment. It is explained why trivial baselines are not redundant measurements and consequently should not be included in GPS networks because they distort the network statistics and make the results look better than they actually are. A consideration of the actual measurements in the process rather than the vectors, is necessary to ensure adequate redundancy in a GPS network.

Overview

Although GPS surveying has been with us for about ten years, there is still disagreement about dependent baselines and redundancy. Confidence in the results of any survey depends upon having a sufficient number of redundant measurements. The concept of redundancy is clear enough. Redundancy is achieved in any system of measurements by having more than the minimum number of measurements to completely define the mathematical model. Redundancy can be achieved with phase-differencing methods by positioning the same point from more than one observing session. One measurement session defines the relative positions of all the points of that session. A second measurement session that includes any point of the first session provides a different set of measurements to the point. The mathematical model defining the point becomes "over determined" and redundancy is achieved. This sounds simple enough, yet there is still much disagreement, in particular about dependent or "trivial" baselines. Trivial baselines appear to be a redundant measurement occurring within one observing session. Some GPS users consider them redundant measurements, while others do not. (Consider Figure 1.)

With receivers simultaneously occupying all of the points of Figure 1, three baseline vectors can be computed. The points appear to be over determined, with one of the baselines providing a redundant measurement. The "redundant" baseline is a dependent vector from a solution of simultaneous equations, so it is called a dependent or "trivial" baseline. The three baselines form a closed figure. If a loop is computed with the vectors, there is a small misclosure that resembles the misclosure of a conventional traverse. Since each baseline is computed separately from the others, there appears to be some redundancy. This "common-sense" approach tells us that there is redundancy; however, a mathematical examination of the process clearly shows that there is none. Why the disparity? Since the problem is a particularly vexing one that fundamentally affects the entire process of surveying with GPS, it warrants a thorough examination.

Part of the difficulty stems from the intangible nature of GPS measurements. What is actually being measured is not obvious. On the other hand, it is much easier to see what is actually being measured with conventional methods. Horizontal and vertical angles are measured. A tape is pulled between two points to measure a slope distance. Even with today's electronic instruments, exactly what is being measured remains clear. Anyone with a high-school background in trigonometry and geometry can intuitively see that these measurements could be used to compute coordinate values. The specifics require a certain amount of training, but it is obvious that it can be done. GPS technology is more nebulous, however. Radio receivers placed on widely separated points receive signals from unseen satellites. Invisible bytes of data are "dumped" into a computer. A set of manufacturer's procedures are followed with the software and out come some apparent

R. Cliff Wilkie is a supervising surveyor for Saudi Aramco. His mailing address is: c/o Saudi Aramco, Box 5713, Dhahran, Saudi Arabia.

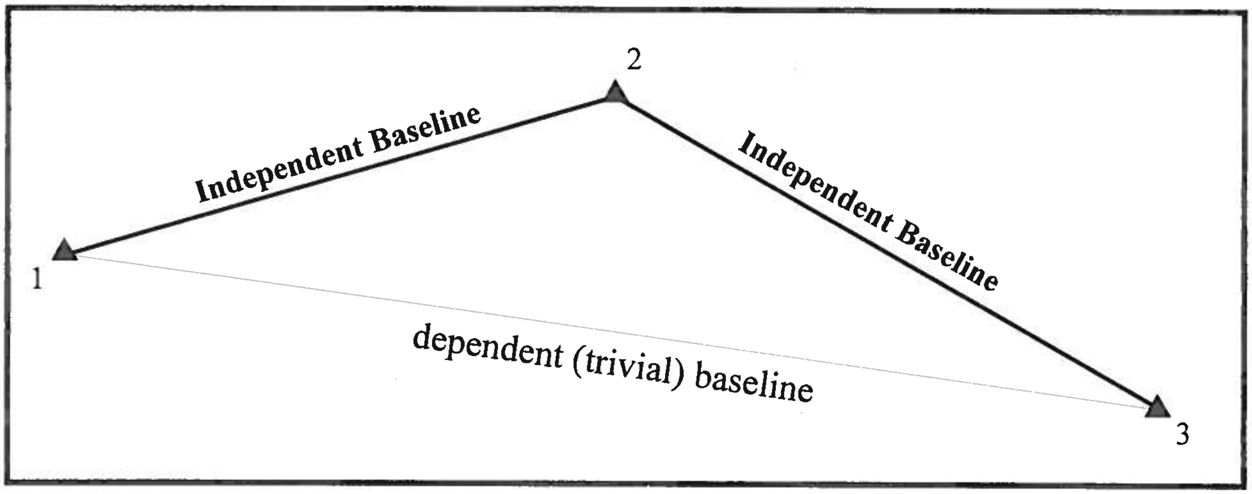

Figure 1. Independent and dependent (trivial) baselines.

"measurements." Some software packages output coordinate differences, others output the coordinates themselves. Either way, the process has a quasi-magical aura to it. There is no obvious source for these answers which are assumed to be measurements. Moreover, if these are not measurements, what are they, and where are the "real" measurements?

In order to clearly understand the relationship between dependent baselines and redundancy, we must first examine what is actually being measured and what is simply being computed. It is useful to view the process in two phases. First, the coordinates of the points are computed by an algebraic process similar to distance-distance intersections. Second, the baselines are computed. The distances involved in the first step are from the different satellites to all the points occupied simultaneously during any session. The baseline vectors are computed as inverses between the points. The fundamental measurements of the process are the distances from the different satellites to all the points occupied simultaneously during the session. Yet even though these are the fundamental measurements of the process, they are not usually entered as observations into a network adjustment. Instead, the coordinate differences between the points (inverses) are considered as observations themselves. So a GPS network is usually put together from computed inverses rather than from actual observations. There is nothing wrong with this practice—it is commonly done as a mathematical convenience in vertical networks. However, it contributes to the confusion about trivial baselines and redundancy. The effects are more obvious when a comparison is made with a similar process which uses conventional surveying methods.

Comparison to Conventional Methods

Suppose we have only two horizontal control points, and want to position two or more new points with only an EDM. We are able to measure only distances, not angles. Any number of new points could be positioned by distance measurements from the two control points (see Figure 2). Each control point would be occupied by the EDM and distances would be measured to all the new points. Distance-distance intersections could be used to generate coordinate values for the new points. The points could be positioned, but there would be no checks of their coordinate values. Inverses could be computed between all the new points, and angles could be computed between the respective lines. A closed loop could even be computed with these angles and distances, but of course there would be no redundancy in the system. If there were a small mistake in all the distances, such as a reflector constant entered with the wrong sign, the mistake would not affect the relative positioning of the new points very much. Since the new points are clustered closely together relative to the distant control points, the effect on their relative positions would be minimal. If it were known, in general, where the points were in relation to the control points, coordinates for as many new points as desired could be computed from only two control points.

The generation of coordinates using GPS phase differencing has some similarities to this process (see Figure 3). The known control points of the previous diagram are replaced by satellites. The satellite coordinates are known from the orbital information broadcast to the user in the ephemeris. The new points of the conventional

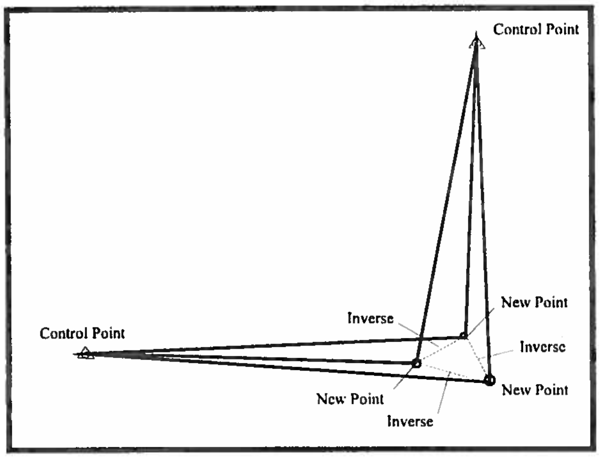

Figure 2. Distance-distance intersection.

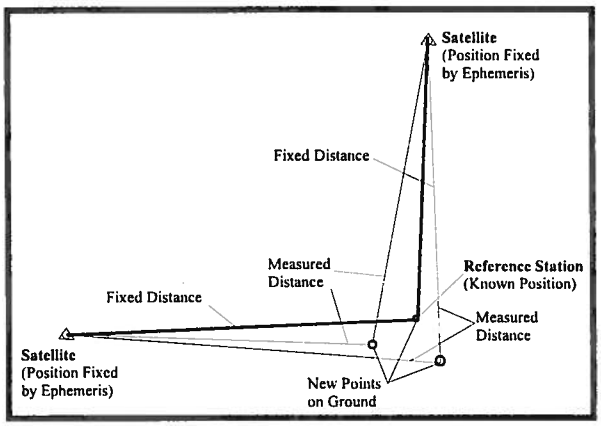

Figure 4. Fixing of reference station.

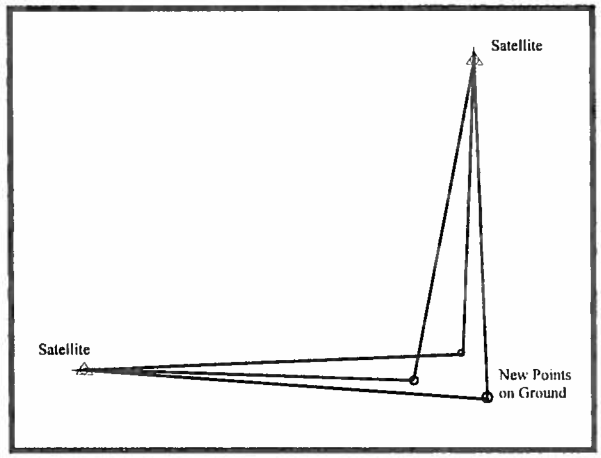

Figure 3. Distance-distance intersections from satellites.

survey shown in Figure 2 are replaced by the new points on the ground in the GPS survey shown in Figure 3. Each one of the new points on the ground is positioned by a receiver capable of measuring distances to the satellites. For various reasons, these distances all contain very similar, but relatively small, errors. Again, the relative positions of the points are not affected very much by these errors. In the hypothetical situation shown in Figure 3, there are only enough satellites to solve a two-dimensional system with no redundancy.

Of course, the process is similar to a distance-distance intersection only in a general way. For a more complete understanding, some differences must to be considered.

Even though the receivers measure the distances from the satellites to all the new points on the ground, not all of these measured distances are used as final distances (see Figure 4). In order to make the mathematical process function properly, one of the new points must be arbitrarily chosen as a reference station and fixed to previously known WGS-84 values. The positions of the satellites are likewise fixed to values computed from the ephemeris data for each epoch of measurement. But even though the phase measurements from the satellites to the reference station are recorded in the receiver at that station, the final distance between the reference station and the satellites is simply the difference between the satellite positions and the nominal value of the reference station. In a sense, therefore, these distances are not measured by carrier phase values. The remaining distances—from the satellites to the other points on the ground—are all "measured" from carrier phase data.

The Phase Difference Solution

In reality, the only thing "measured" by the receiver is the partial phase of a cycle. In other words, when a satellite first comes into the view of the receiver, the receiver is capable of measuring the portion of one whole cycle which is present at the receiver's antenna at the moment of measurement.

This part of one whole cycle is measured and stored as a real number less than one. For each subsequent measurement interval (epoch), the receiver is capable of both counting how many whole cycles have elapsed since the last measurement and measuring whatever portion of a cycle is present at the moment of the subsequent measurement. So, only the partial phase of the first cycle is measured for the first epoch, rather than the entire distance from the satellite to the ground. In order to obtain the entire distance, something has to be done to derive the whole number of cycles from the receiver to the satellite at the first moment of measurement.

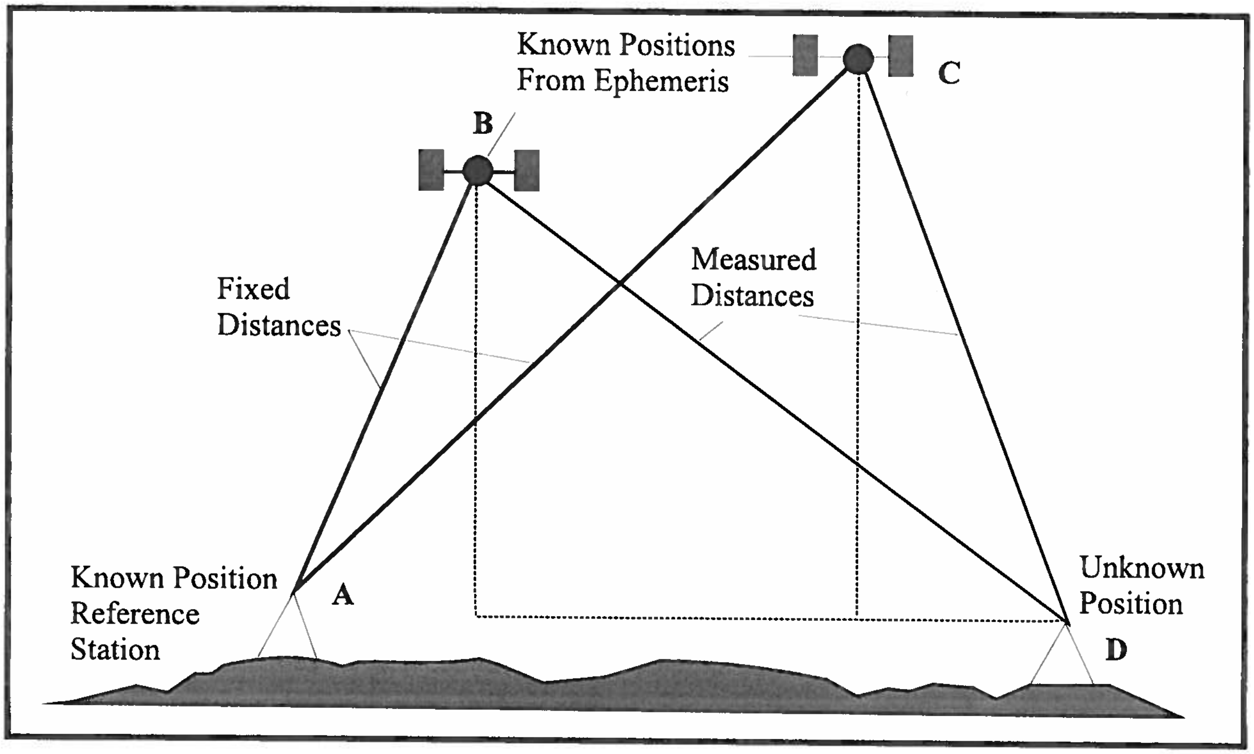

Figure 5. Double difference.

A variety of intricate methods have been developed to do this. Much of GPS software development focuses on this area and tremendous strides have been made. Until recently, the whole number of cycles to each satellite was treated first as a parameter, along with others, in a general case least squares adjustment where both parameters and observations are mixed in condition equations. (This general or unified approach to least squares adjustments is outlined in Mikhail 1976 and Mikhail and Gracie 1981.) If the parameter solutions were close to integer values, the parameters were changed to constants by simply rounding off to the nearest integer value and re-solving the system with that many fewer parameters (Leick 1990). More recent procedures use a sophisticated series of differencing techniques to approximate a final solution, then test all the possible integer combinations within a very limited range of possibilities. If a combination can be found with significantly smaller error statistics, it is considered correct (Allison 1991, Frei 1990, Talbot 1992).

Once this value is computed, it is tacked onto all the phase measurements for every epoch, and the entire distance from the satellite to receiver becomes a measurement, at least in effect. These distances are the fundamental measurements of GPS positioning. They are analogous to the fundamental measurements of conventional surveying, horizontal angles, vertical angles, and slope distances. Once they are derived, these distances are differenced against the fixed distances from the satellites to the reference station to compute the coordinate values of the other points on the ground.

In practice, the entire distance from the receiver to each satellite is never computed and expressed as a measurement. It is only the differences between these distances and the fixed distance to the reference station that are derived and shown to the user via the software.

Now the difference between the process and a distance-distance intersection becomes more apparent. Some of the distances are fixed and the others are considered as measured values. These quantities become tools for determining the whole number of cycles from the satellites to the ground points. Once the whole number of cycles is determined for each distance, it is added to the phase measurements and the entire quantity is used as a measurement in determining the coordinates for the unknown points.

Various differencing methods are used before a final solution is reached, but the double difference solution shown in Figure 5 illustrates the concepts well. The solution is based fundamentally in a solution of right triangles. The measured distances of the diagram are considered the known hypotenuses of right triangles. The coordinates of the unknown position are algebraically solved with simultaneous equations modeling the adjacent sides of the triangles as coordinate differences between the satellite coordinates and the coordinates of the unknown point. The satellite coordinates are known quantities. Coordinates of the receiver at the unknown point are the unknown quantities in the diagram. These are what actually "drop out" of the solution.

In the situation shown in Figure 5, two satellites and two measured distances are enough to solve for two-dimensional coordinates. Since GPS positioning is done in three dimensions, another satellite is needed for three-dimensional coordinates. And since the whole process occurs over a changing time interval, another satellite is needed for time. Thus a minimum of four satellites are necessary. The actual process is a rather complex solution of simultaneous equations. There is a massive amount of data, so matrix methods are used. Regardless of the complexity of the actual solution, the process is not, in essence, too far removed from the diagram of Figure 5. Looking at the diagram, it is intuitively clear that, regardless of the specifics of the algebra, enough information is available to determine coordinates for the unknown point D with the two measured distances BD and CD. This is actually a simultaneous mathematical process, but treating the process as a sequence of mathematical events makes it easier to understand the concepts.

Effects of Errors in the Fixed Distances

The coordinates of any unknown points each contain absolute positional errors resulting from the errors in the fixed distances from the satellites to the reference station. These errors arise from a combination of errors both in the position of the reference station and in the satellite positions. The effects of these errors in the absolute positions of the unknown points have very little effect on the relative positional accuracy of the points. This can be visualized by a re-examination of Figure 4.

Figure 4 is a somewhat more realistic representation of a typical geometry of points and satellites than is Figure 5. The errors in the fixed distances are normally small in relation to the total distance to the satellites. Further, the unknown points are clustered together in an area which is small when compared with the area occupied by satellites spread over the sky at distances greater than 20,000 km from the unknown points. The accuracy of the unknown points with respect to each other is more of a function of the precision of the phase measurements rather than the error in the fixed distance. This relative positioning accuracy deteriorates as the unknown points become further apart and/or the satellites become closer together.

Session Results Considered as Points rather than Vectors

The above method is appropriately called relative positioning. However, to clearly evaluate the effects of trivial baselines on redundancy, the positions of the points in relation to the satellites needs to be considered, rather than the positions of the points in relation to each other. This shifts the focus to the measurements and away from the three-dimensional vectors that directly express the relative positions of the points. It is easier to visualize this if the results of a session are thought of only as the coordinates that "drop out" of the solution, rather than as vectors between the points. The results of each simultaneous session are viewed as a set of three-dimensional coordinates floating in space, rather than as interconnected vectors. Each additional session produces another set of coordinates floating in space in relation to a different satellite geometry. If the same point is positioned in more than one session, more than one set of coordinates will exist for it—a redundant session has been measured to the point.

Some contemporary software outputs only the coordinates of the new points to the user (Wild, Leitz 1989). Others take the process an additional step, computing coordinate differences between the new points and outputting these data to the user for input as vectors into a network adjustment (Trimble Navigation 1994). This is a critical juncture where the user can easily mislead himself into thinking of the vectors between points as the measurements themselves. Since they are simply inverses between coordinate values, their effect on redundancy has to be considered differently from that of observations.

External and Internal Redundancy

Unfortunately the issues regarding redundancy are more complex. A further distinction has to be made between the "external" redundancy achieved from different sessions and the

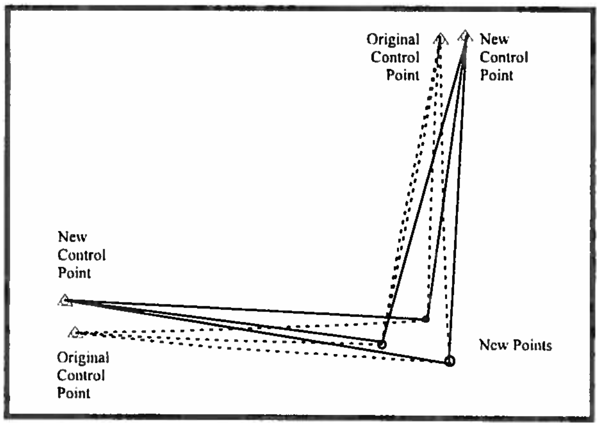

Figure 6. Distance-distance intersections with new control points.

"internal" redundancy of any one session (Pfeifer 1982). There is a good deal of internal measurement redundancy in any one session, but significant errors and potential sources of mistakes external to one session remain which must be resolved by a different session. Internal redundancy is illustrated by Figures 2, 6, and 7. It is apparent from Figure 2 that even though many points can be positioned from the two control points, there is no redundancy in the system. One way to achieve redundancy would be to position the points from one or more different control points (see Figure 6).

With the distances from the new control points, we could compute another set of coordinates for the new points. Since the new control points are close to the first set of control points, the coordinate values will be only slightly different from the first values. This is also done in GPS computations except that, instead of different control points, we use the same satellites at a different time, so their positions in orbit are different (See Figure 7).

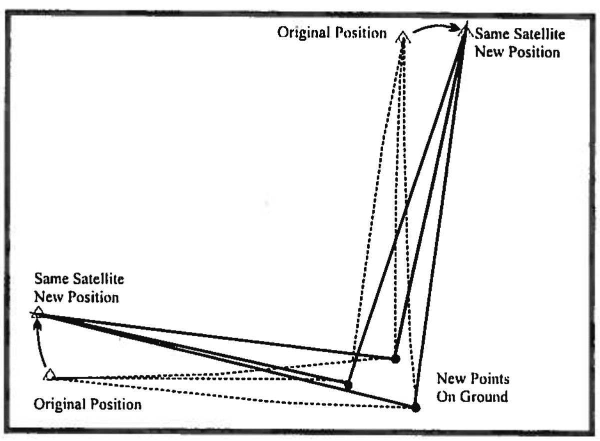

Figure 7. Distance-distance intersections from new satellite positions.

This provides some redundancy within, or internal to, one particular session of measurements.

The further the satellites move from their initial positions, the more accurate the redundancy check. As observations are taken to moving satellites, each measurement epoch provides a unique set of measurements that provide redundancy within a particular session. Each satellite above the minimum of four also provides redundancy. But with typical epoch intervals of less than one minute, the satellites have not moved very much in their orbits from one epoch to the next. As a result, measurements from successive epochs contain virtually no difference in their errors, and consequently produce only very slight differences in coordinate values of the points. So, even though a vector computed between two stations by measurements from one session may contain several hundred measurements, it is usually reduced to only one vector and considered as only one measurement in a network. This is done even though there are many measurements within that one session and there is quite a bit of internal redundancy within that one set of measurements.

This is somewhat similar to what is done in conventional surveying when a complete set of perhaps as many as 16 positions of an angle is reduced to one measured angle and entered into an adjustment as one observation, not 16. There is obviously some internal redundancy. The angle has been measured 16 times when only one angle is sufficient to define the mathematical model. But there is no redundancy that models external factors such as instrument and target setup errors, or perhaps, more importantly, instrument and target setup mistakes. There is nothing to indicate that the targets were over the correct point or that a tribrach was not leveled, or was badly out of adjustment. There is some canceling of errors resulting from reading two faces, rotating the circle, and so on, but, essentially, the process simply refines the precision of the resultant angle and produces a tighter set of statistics. However, conventional methods also provide a means of dealing with external factors since each station occupied by a theodolite is also a target station for subsequent sets of angles measured from backsight and foresight stations.

In a similar fashion and for the same reasons, each point in a GPS network should be occupied more than once. There is some "internal" redundancy in only one session since there are more than the minimum number of satellites and epochs. But basically the process of observing for many epochs simply makes the integer ambiguities (parameters) converge closer to integer values and produces a tighter covariance matrix. So, normally, a complete set of all the measurements of any one session is considered as only one set of measurements for computing coordinate values of the new points being positioned. There may a great deal of internal redundancy from that one session, but there is still nothing to model receiver-setup error or mistakes. Different sessions, especially if done on different days and with different satellites are what are needed to achieve external redundancy in the network. They provide a completely different satellite geometry and completely different sets of measurements from the satellites to the new points.

Loop Misclosures with Trivial Baselines

Another issue that needs to be considered is the closed loop that can be formed by a trivial baseline. It is apparent from looking at Figure 2 that, if inverses are computed between all three of the new points and considered as measurements for a traverse between the points, the traverse will close perfectly. The same thing happens in a GPS session. If all three vectors are used in a loop or network, they will close almost perfectly and for the same reasons. The baselines are simply inverses between coordinate values all generated from the same set of distances to the satellites. With GPS measurement the loop will not close perfectly, although it will usually be very close. To most surveyors, the small misclosure appears equivalent to a traverse misclosure. It is not the same thing.

In order to fully understand why the closing vector is not a redundant measurement, it is important to understand the reasons behind the small misclosure. Various factors contribute. Part of the misclosure stems from computing vectors of a session separately rather than computing all the vectors of a session together. When vectors are computed separately, the fixed distances are used again for each vector. When all the vectors are computed together, the fixed distances are only used once. This is more appropriate mathematically, but not quite as practical if vectors are to be used in a network. So, many users compute vectors separately. The small differences that result from doing this contribute towards a misclosure. Other factors enter the picture. Sometimes the data from different vectors of a given session are not exactly simultaneous. Most of the observations occur over the same period, but one vector may have a few epochs more or less than another, so the vectors vary slightly from what would be produced from completely simultaneous data. Another contributing factor relates to the reference station. Since each of the vectors is usually computed separately, different reference stations are sometimes used. Since these may be from independently derived pseudorange positions, each position will have completely different amounts of error. If the vectors are not computed from a common reference station, the errors in the pseudorange positions will contribute something towards the misclosure. Occasionally, one vector from an otherwise satisfactory session will produce a large misclosure. Usually this is because the full number of cycles for that particular data set are incorrectly solved and, accordingly, the software produces a completely spurious vector. This is similar to what is called in conventional surveying a blunder or mistake rather than an error in a strict mathematical sense. Suffice it to say that if the GPS vectors of a particular session were all computed in relation to one reference station, and with simultaneously observed data, then a loop computed between the new stations of the particular session would close perfectly. Thus, even though the small misclosure may look like a conventional traverse misclosure, the differences in closing coordinate values come from different sources.

Why the Trivial Baseline Is Not a Redundant Measurement

By looking at the actual measured values to the satellites rather than at the vectors between the points, it can be seen why a trivial vector does not provide redundancy. With receivers simultaneously occupying three points, three unique sets of measurements are collected from the satellites to the points. These three sets of measurements can be used in differing combinations to compute a total of three inverses between the points. If one station is arbitrarily fixed to known values, then all the distances between that station and the satellites become fixed distances. The remaining distances measured at the two other points are used to compute an inverse, first between the fixed station and one point, then between the fixed station and the remaining point. Now all the actual measurements between the satellites and the ground points have been "used up." A third vector can be computed by fixing another station as a reference and using previously used distances to compute the final vector. But the third vector is computed from distances that have already been

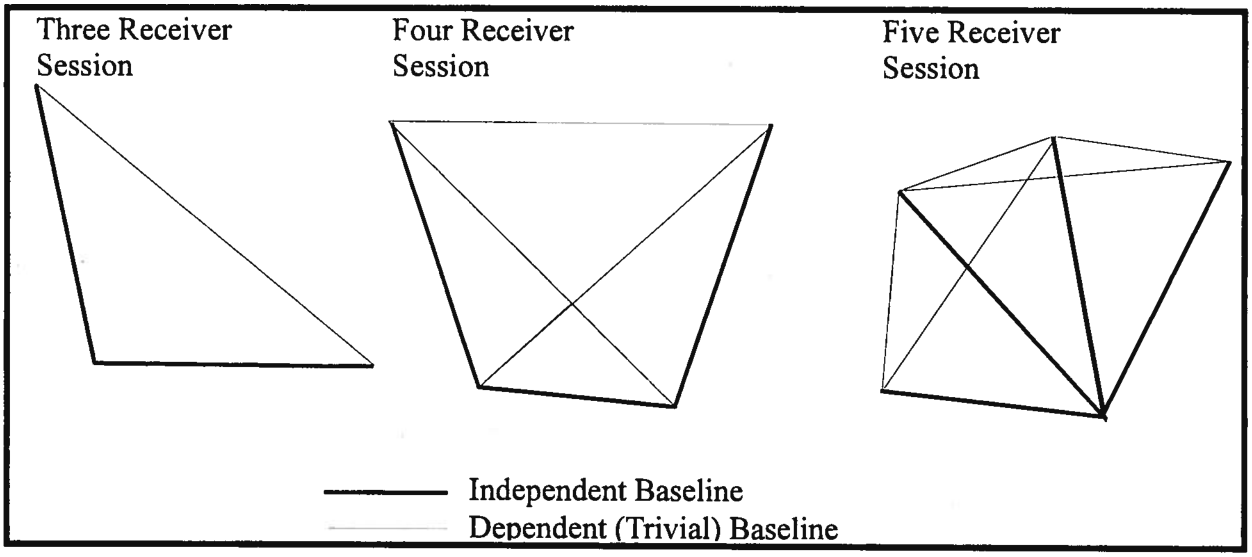

Figure 8. Dependent and independent baselines with differing numbers of receivers.

used to compute the other two vectors, so there is no redundancy from the third vector. It may be convenient to have all three computed in order to have all the possible choices available from the session, but one of the three should ultimately be discarded. Including it in an adjustment is tantamount to entering the same measurement again. It would be somewhat like entering the same angle or distance twice in an adjustment of a set of conventional observations.

Statistical Effects of Including a Trivial Baseline

Since the measured distances to the satellites are not directly entered twice, entering a trivial baseline is not quite the same as doubling a measurement. It does vary the relative coordinate values of the points of a given session because the adjustment process "thinks" the dependent vector is a redundant measurement. It also gives more weight to the vectors of a given session in relation to the other sessions of the network. The final baseline precisions are made to look better than they otherwise would by a smaller variance factor. Most network software contains a feature for multiplying a variance factor back into the a posteriori cofactor matrix. This normally should be done since GPS measurements usually do not have realistic a priori weights. Without realistic a priori weights, the variance factor becomes the only way of relating the statistical analysis of the network to the actual residuals of the observations. The variance factor is a fraction which could be thought of as the variance of the entire network. The numerator is a weighted sum of the residuals squared, and the denominator is the number of degrees of freedom in the network. Adding dependent vectors distorts the variance factor. Each dependent vector adds three degrees of freedom to the network and consequently adds the number three to the denominator of the variance factor. Meanwhile, the numerator is only increased by the square of the weighted residuals of the vector. This is normally a very small number, less than one. So the numerator is only increased slightly, the denominator is increased by three, and the resulting quotient becomes a smaller value. As more dependent vectors are added, the quotient becomes even smaller. This is then multiplied into the a posteriori cofactor matrix to convert it into a covariance matrix which is then used to compute baseline precisions (Mikhail and Gracie 1981). With a smaller variance factor, the values in the covariance matrix become smaller. The smaller variances make the baseline ppm values look better than they otherwise would. The effects are relatively small with only three receivers but increase greatly with additional receivers.

The ratio of dependent to independent vectors increases rapidly with each additional receiver. With three receivers, there is only one trivial baseline with two independent baselines. With four receivers, three independent vectors can be computed and three trivial vectors. With five receivers, four independent vectors result and six trivial can be produced. There is always one less usable vector from any session than the number of receivers. With five receivers there are four usable vectors. With ten receivers, there are nine vectors. We can pick and choose between all the vectors that can

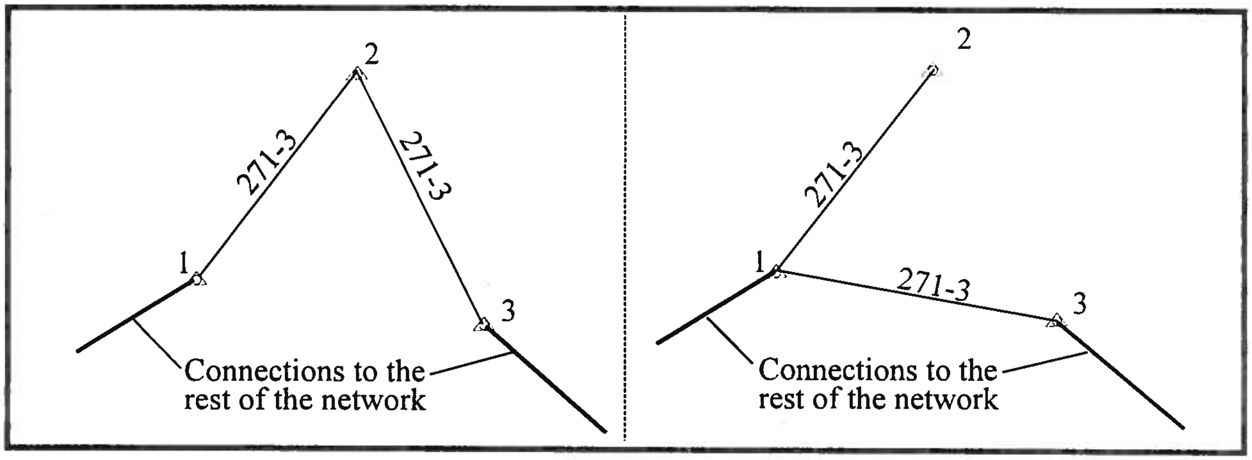

Figure 9. Single session connection of a point to a network.

be computed from any one session to decide which ones we want to use as independent vectors, but the total number cannot be more than one less than the total number of receivers used in the session, and none of them can form a closed loop.

Figure 8 contains diagrams of observing sessions with differing numbers of receivers.

Choice of Independent Vectors and Effects on Point Redundancy

Putting together a network of three-dimensional coordinate differences requires a judicious evaluation of which vectors are considered independent and which dependent, or some of the points will be inadvertently positioned without a redundant session. Surprisingly, it makes no difference which vectors of a particular session are considered independent and which dependent. This is often misunderstood, and favoring certain choices over others can give an appearance of redundancy where none is actually present. In the final analysis, every individual point has to be positioned by vectors from at least two separate sessions, or there is no external redundancy for that particular point.

Consider the example illustrated in Figure 9: The three points 1, 2 and 3 were all positioned simultaneously during the third observing session of the 271st day of the year (271-3). Any two baselines between the three points can be considered as independent. Most conventional surveyors would prefer the choice on the left because it appears to be more like a closed traverse through point 2. But, in practice, positioning point two by drawing the connecting vectors as shown on the left is no better than connecting them as shown on the right. Either way, there is no redundant session at point two. The situation is equivalent to conventional positioning of a geodetic network point by a sideshot rather than including it within a network or as a traverse point. The most significant danger of this method is the possibility of an undetected blunder in measuring the antenna height or of not centering carefully over the point. Further, any biases caused by ionospheric variations, orbital errors or anything else, would go undetected. In order to achieve the necessary redundancy, point 2 must be connected to the network by vectors from at least two different sessions. (See Figure 10.)

It is usually easier to visualize these concepts by considering any points positioned from a given session as coordinates floating in space rather than as connected vectors between the points. It is then more obvious which points are positioned with a redundant session and which are not. Regardless of the intellectual approach used, each point in a network, whether an interior point or a point on the outside border of the network, must be positioned from at least two sessions, or there is not adequate redundancy of the coordinate values of that point. If thought of in terms of baseline vectors, then each point must be connected by vectors from at least two different sessions.

Closing Commentary

GPS methods are deceptively similar to conventional methods, and the similarities cause confusion about how to achieve redundancy with GPS. Since GPS vectors are considered observations in network adjustments, it is easy to jump to the conclusion that they are in fact, observations. So computing a dependent vector as a "closing line"

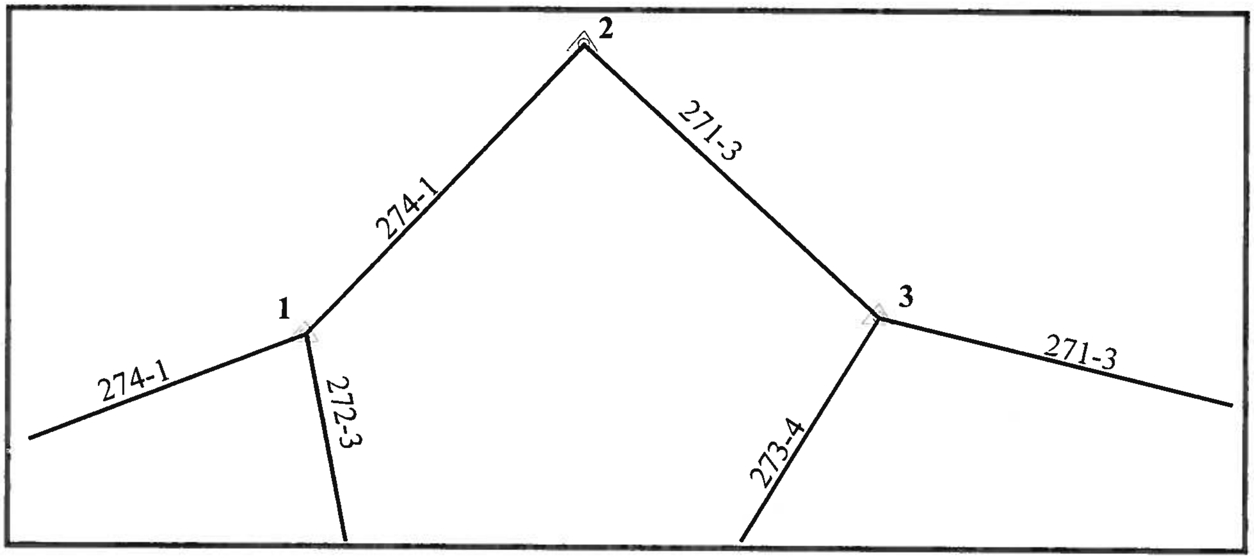

Figure 10. Multiple session connection of a point to a network.

is an appealing way to attempt to achieve redundancy. However, in order to achieve full redundancy with GPS methods, the actual observations have to be considered, not the computed vectors. The computation of a dependent vector does not involve any more observed distances to the satellites than does the computation of only the independent vectors, so no redundancy is effected by dependent vectors.

For a long time surveyors have used conventional methods to produce coordinate values for unknown points. Naturally, when a new technology comes along which enables them to produce coordinate values more cheaply and easily, they are going to invest in it and put it to good use. GPS seems to be simply another marvel from the same fountain of technology that has produced electronic handheld calculators, computers, electronic distance-measuring equipment, electronic theodolites, and so forth. However, the mathematical underpinnings of GPS positioning are fundamentally different from those of conventional surveying.

Changing from a total station and conventional methods to GPS methods is not simply another step in the evolution of surveying technology. Some analogies can be made which bridge the gap between the two areas; however, the mathematics of GPS positioning need to be properly understood before the method can be fully exploited and safely applied. Mathematical intuition is a function of mathematical background. The intuitive leap that enables a person to "see" how to compute coordinates from conventional measurements will not necessarily serve to make the same leap with GPS methods.

ACKNOWLEDGMENTS

The author would like to thank Subramaniam Ganeshu for being a patient sounding board for many of the ideas in this paper, Bob Mandzi for his critical commentary, and especially the four anonymous reviewers as well as the editor of this journal for the candor and thoroughness of their commentaries.

REFERENCES

Allison, T. 1991. Multi-observable processing techniques for precise relative positioning. *GPS-91 Proceedings, ION 4th International Technical Meeting*.

Frei, E., and G. Beutler. 1990. Rapid static positioning based on the fast ambiguity resolution approach "FARA": theory and first results. *Manuscripta Geodaetica* 15.

Leick, A. 1990. *GPS satellite surveying*. New York: Wiley.

Mikhail, M. E. 1976. *Observations and least squares*. Lanham, Md., New York, and London: University Press of America.

Mikhail M. E., and G. Gracie. 1981. *Analysis and adjustment of survey measurements*. New York: Van Nostrand Reinhold.

Pfeifer, L. 1982. Input formats and specifications of the National Geodetic Survey Data Base, volume 1. In *Horizontal control data, NOAA manual NOS NGS 2.: External and internal errors*. Silver Spring, Md.: NOAA, 2-16

Talbot, N. C. 1992. *Recent advances in GPS surveying*. Sunnyvale, Calif.:Trimble Navigation.

Trimble Navigation Ltd. 1994. *GPSurvey software user's manuals*. Sunnyvale, Calif.: Trimble.

Wild, Leitz Ltd. 1989 *POPS software manual*. Norcross, Ga.: Wild, Leitz.

FURTHER READING

GeoComp, Inc. 1989-1991. *Turbonet software manual*.

GEOsurv Inc. 1994. *Geolab software manuals*. Ottawa, Canada: GEOsurv.

Trimble Navigation Ltd. 1991 *TRIMVEC-PLUS, GPS survey software manual*. Sunnyvale, Calif.: Trimble.

Vincenty, T.1986. *Fillnet Software Manual*. ■